

CLAIM:

1. In a RAID data storage system, a method comprising:
receiving a first request to read data of a stripe unit B_x of a stripe, wherein the first request is received from a computer system in data communication with the RAID data storage system;
returning data of stripe unit B_x to the computer system in response to receiving the first request;
receiving a second request to read data of stripe unit B_x , wherein the second request is received from the computer system;
generating new data for stripe unit B_x in response to receiving the second request;
returning the new data to the computer system.
2. The method of claim 1 wherein the new data is generated as a function of error correction data and data of stripe units of the stripe other than stripe unit B_x .
3. The method of claim 2 wherein the new data for stripe unit B_x is generated only if the second request is received within a predetermined amount of time after receiving the first request.
4. The method of claim 1 further comprising overwriting data of stripe unit B_x with the new data unless a third request to read data of stripe unit B_x is received from the computer system within a predetermined amount of time after the second request is received.
5. The method of claim 1 further comprising comparing data of stripe unit B_x with the new data.
6. The method of claim 5 further comprising notifying the computer system of data inconsistency if the comparison shows the data in stripe unit B_x is different from the new data.
7. The method of claim 5 further comprising failing the second request for data in stripe unit B_x .

8. The method of claim 7 wherein the failure is communicated to the computer system as an error that indicates a data corruption condition.
9. The method of claim 5 further comprising determining that the parity data for the stripe containing stripe unit B_x is corrupted and repairing it.
10. The method of claim 9 wherein new parity information for the stripe unit B_x is recalculated from data available in all stripe units of the stripe containing stripe unit B_x .
11. The method of claim 9 wherein the new parity information is written back to a parity block of the stripe containing stripe unit B_x .
12. The method of claim 5 further comprising determining that data stored in stripe unit B_x is corrupted and repairing it.
13. The method of claim 12 wherein the new data for stripe unit B_x is recalculated from all other stripe units and parity of the stripe containing stripe unit B_x .
14. The method of claim 12 wherein the new data for stripe unit B_x is written to stripe unit B_x .
15. In a RAID data storage system that receives a request to read data stored in a stripe unit B_x , wherein the request is received from a computer system in data communication with the RAID data storage system, a method comprising:
 - comparing an identification for stripe unit B_x with identifications for stripe units stored in a table in memory;
 - if the identification for stripe unit B_x does not compare equally with one of the identifications stored in the table:
 - storing the identification for stripe unit B_x in the table, and;
 - returning data of the stripe unit B_x to the computer system;
 - if the identification for stripe unit B_x does compare equally with one of the identifications stored in the table:
 - generating new data for stripe unit B_x , and;
 - returning the new data to the computer system.

16. The method of claim 15:

wherein stripe unit B_x is one of a plurality of stripe units of a stripe;

wherein the new data is generated as a function of (1) error correcting data for the stripe and (2) data of the plurality of stripe units other than stripe unit B_x .

17. The method of claim 15 further comprising:

storing a copy of the new data in memory;

overwriting data of stripe unit B_x with the copy of the new data stored in memory if the computer system determines that the new data returned to it is valid.

18. The method of claim 15 further comprising:

storing a copy of the new data in memory;

overwriting data of stripe unit B_x with the copy of the new data stored in memory unless a new request is received from the computer system to read data stored in stripe unit B_x within a predetermined amount of time after generating the new data.

19. In a RAID data storage system, a method comprising:

receiving a first request to read data of a plurality of stripe units, wherein the first request is received from a computer system in data communication with the RAID data storage system;

returning data of the plurality of stripe units to the computer system in response to receiving the first request;

receiving a second request to read data of the plurality of stripe units, wherein the second request is received from the computer system;

generating new data for a first stripe unit of the plurality of stripe units in response to receiving the second request;

returning the new data to the computer system along with data of the plurality of stripe units other than the first stripe unit.

20. The method of claim 19 wherein the new data is generated as a function of error correction data and data of stripe units other than stripe unit B_x .

21. The method of claim 19 wherein the new data for stripe unit B_x is generated only if the second request is received within a predetermined amount of time after receiving the first request.

22. The method of claim 19 wherein the plurality of stripe units consist of stripe units from first and second separate stripes.

23. A computer readable medium storing instructions executable by a first computer system in a RAID data storage system, wherein the RAID data storage system comprises a stripe, wherein the stripe comprises a plurality of stripe units including a stripe unit B_x , wherein the first computer system performs a method in response to executing instructions stored on the computer readable medium, the method comprising:

returning data of stripe unit B_x to a second computer system coupled to the first computer system, wherein the data of stripe unit B_x is returned in response to receiving a first request to read data of the stripe unit B_x , wherein the first request is received from the second computer system;

receiving a second request to read data of stripe unit B_x , wherein the second request is received from the second computer system;

generating new data for stripe unit B_x in response to receiving the second request;

returning the new data to the second computer system.

24. The computer readable medium of claim 23, wherein the new data is generated as a function of error correction data and data of stripe units of the stripe other than stripe unit B_x .

25. The computer readable medium of claim 24, wherein the new data for stripe unit B_x is generated only if the second request is received within a predetermined amount of time after receiving the first request.

26. The computer readable medium of claim 23, wherein the method further comprises overwriting data of stripe unit B_x with the new data unless a third request to read data of stripe unit B_x is received from the second computer system within a predetermined amount of time after the second request is received.

27. A computer readable medium storing instructions executable by a first computer system in a RAID data storage system, wherein the RAID data storage system comprises a stripe, wherein the stripe comprises a plurality of stripe units including a stripe unit B_x , wherein the first computer system performs a method in response to executing instructions stored on the computer readable medium, the method comprising:

- comparing an identification for stripe unit B_x with identifications for stripe units stored in a table in memory in response to the first computer system receiving a request to read data stored in stripe unit B_x , wherein the request is received from a second computer system;

- if the identification for stripe unit B_x does not compare equally with one of the identifications stored in the table:

- storing the identification for stripe unit B_x in the table, and;

- returning data of the stripe unit B_x to the second computer system;

- if the identification for stripe unit B_x does compare equally with one of the identifications stored in the table:

- generating new data for stripe unit B_x , and;

- returning the new data to the second computer system.

28. The computer readable medium of claim 27:

- wherein stripe unit B_x is one of a plurality of stripe units of a stripe;

- wherein the new data is generated as a function of (1) error correcting data for the stripe and (2) data of the plurality of stripe units other than stripe unit B_x .

29. The computer readable medium of claim 28, wherein the error correcting data comprises parity data for the stripe.

30. The computer readable medium of claim 27, wherein the method further comprises:

storing a copy of the new data in memory;
overwriting data of stripe unit B_x with the copy of the new data stored in memory if the second computer system determines that the new data returned to it is valid.

31. The computer readable medium of claim 27, wherein the method further comprises:

storing a copy of the new data in memory;
overwriting data of stripe unit B_x with the copy of the new data stored in memory unless a new request is received from the computer system to read data stored in stripe unit B_x within a predetermined amount of time after generating the new data.

32. The computer readable medium of claim 27, wherein the method further comprises removing the identification for stripe unit B_x from the table unless a new request is received from the second computer system to read data stored in stripe unit B_x within a predetermined amount of time after the identification is stored in the table.

33. A computer readable medium storing instructions executable by a first computer system in a RAID data storage system, wherein the first computer system performs a method in response to executing instructions stored on the computer readable medium, the method comprising::

returning data of a plurality of stripe units to a second computer system in response to receiving a first request to read data of the plurality of stripe units, wherein the first request is received from the second computer system, wherein the second computer system is in data communication with the first computer system;
receiving a second request to read data of the plurality of stripe units, wherein the second request is received from the second computer system;
generating new data for a first stripe unit of the plurality of stripe units in response to receiving the second request;
returning the new data to the second computer system along with data of the plurality of stripe units other than the first stripe unit.

34. The computer readable medium of claim 33, wherein the new data is generated as a function of error correction data and data of stripe units other than stripe unit B_x .

35. The computer readable medium of claim 34, wherein the new data for stripe unit B_x is generated only if the second request is received within a predetermined amount of time after receiving the first request.

36. The computer readable medium of claim 33, wherein the method further comprises overwriting data of stripe unit B_x with the new data unless a third request to read data of stripe unit B_x is received from the second computer system within a predetermined amount of time after the second request is received.

37. The computer readable medium of claim 33, wherein the plurality of stripe units consist of stripe units from first and second separate stripes.

38. In a RAID data storage system, an apparatus comprising:
means for receiving a first request to read data of a stripe unit B_x of a stripe, wherein the first request is received from a computer system in data communication with the RAID data storage system;
means for returning data of stripe unit B_x to the computer system in response to receiving the first request;
means for receiving a second request to read data of stripe unit B_x , wherein the second request is received from the computer system;
means for generating new data for stripe unit B_x in response to receiving the second request;
means for returning the new data to the computer system.